Includes Real-World Scenarios, Hands-On Exercises, and Leading-Edge Exam Prep Software Featuring:

* Custom Test Engine
* Hundreds of Sample Questions
* Electronic Flashcards
* Entire Book in PDF

MCTS

Microsoft® Exchange Server 2007 Configuration

STUDY GUIDE

Second Edition   Exam 70-236   **Joel Stidley**

SYBEX   SERIOUS SKILLS.

Chapter

1

Preparing for the Exchange Installation

MICROSOFT EXAM OBJECTIVES COVERED IN THIS CHAPTER:

✓ **Installing and Configuring Microsoft Exchange Servers**

✓ **Prepare the infrastructure for Exchange installation.**

If it isn't already clear from the title, the primary goal of this book is to prepare you to pass the 70-236 exam. This being the case, we'll spend most of our time together ensuring that you acquire the required knowledge and skills to help you achieve that goal. As someone who has a great deal of passion for messaging, I also hope not only to help you to be successful on the exam, but also to be successful as a messaging professional.

In this chapter, we will start with some of the basics of Exchange and Active Directory. Later on in this chapter, we will look at what's new in Exchange Server 2007. This should help those who have used previous versions of Exchange Server ramp up on key new features. As part of that discussion, we will also cover what is no longer included in or supported by Exchange Server 2007. In later chapters in the book, you'll dig deeper into key concepts and core skills that will prove to be important in your day-to-day administration of Exchange Server 2007 and, of course, important to you on exam day.

This chapter provides you with a good conceptual background of the topics covered in the remainder of the book. Specifically, we will address the following:

- Overview of Exchange Server

- What is new and what has been removed in Exchange Server 2007

- Active Directory and its integration with Exchange Server 2007

What Is Exchange Server 2007?

The fact that you are reading this book means that you probably have a basic understanding of what Exchange Server is about and what it is used for. To set the stage for the remainder of the book, let's review a little bit of the history of Exchange Server.

Exchange was introduced as a Microsoft product in 1996, as the eventual full replacement of Microsoft Mail. Exchange 4.0, as it was called, was an X.400-based messaging system that introduced us to features such as public folders and calendaring functions. It didn't fully embrace the Internet until 1997, however, when Exchange 5.5 was released and we were able to use Outlook Web Access and send and receive SMTP email.

The next version, Exchange 2000, maintained most of the features from the previous releases except for the built-in directory service, which contained configuration and

information about mailboxes and recipients. What filled this void? A descendant of the Exchange 4.0 directory was reworked into a more powerful version that was then built into Windows 2000 and called Active Directory or, as the service is called in Windows Server 2008, Active Directory Domain Services (AD DS). Although this was a painful transition for many organizations due to the complexity of the migration, it turned out to be the right direction for Exchange as a product. This change allowed Exchange to become more flexible and more scalable.

Exchange 2003 was released three years later with improved scalability, stability, and mobility. Features like RPC over HTTP, Recovery storage groups, Exchange ActiveSync, and the Exchange migration tools made it one of the most compelling yet easiest versions to deploy to date.

Even with the improvements found in Exchange 2003, there were areas that needed work. Exchange 2007 was released as a 64-bit-only application, requiring the use of 64-bit-capable hardware with a 64-bit edition of Windows Server. It also introduced the concept of server roles, allowing specific features of Exchange to be installed on separate servers. Two of these five new roles are the Unified Messaging role, which provides for integrating voicemail and fax features, and the Edge Transport role, which is designed as an Internet-facing mail-processing engine.

What is Exchange Server 2007? Simply put, it is an enterprise-class messaging system that provides the best-in-class email delivery, unified messaging, and electronic calendaring functionality.

Active Directory for Exchange Server 2007

As briefly discussed in the last section, *Active Directory* is one of the most important components of Exchange Server. Although a full discussion of Active Directory is outside the scope of this book, the nature of Exchange Server's tight integration with Active Directory warrants a brief discussion of the technology and an examination of how it affects the Exchange messaging environment.

Active Directory

To understand Active Directory, it is first necessary to understand what a directory is. Put simply, a *directory* contains a hierarchy that stores information about objects in a system. This is similar to how a phone directory stores information about a person, their phone number, and their home address.

> **NOTE** Windows Server 2008 introduced Active Directory Domain Services (AD DS). This is essentially a rebranding of Active Directory to describe the feature better and to be able to incorporate related products with Active Directory branding in Windows Server 2008. These products include Active Directory Certificate Services (AD CS), Active Directory Lightweight Directory Services (AD LDS), Active Directory Federation Services (AD FS), and Active Directory Rights Management Services (AD RMS). Since Active Directory for Exchange Server 2007 works with both Windows Server 2003 and 2008 and the core functionality is the same, this book will generically refer to it either as Active Directory or AD.

A *directory service* is the service that manages the directory and makes it available to users on the network. Active Directory stores information about objects on a Windows Server network and makes it easy for administrators and users to find and use it. Active Directory uses a structured data store as the basis for a hierarchical organization of directory information.

You can use Active Directory to design a directory structure tailored to your organization's administrative needs. For example, you can scale Active Directory from a single computer network all the way to many networks. Active Directory can include every object, server, and domain in a network.

What makes Active Directory so powerful and so scalable is that it separates the logical structure of the Windows Server domain hierarchy from the physical structure of the network.

Logical Components

In Exchange 5.5 Server and prior versions, resources were organized separately in Windows and in Exchange. Now the organization you set up in Active Directory and the organization you set up in Exchange Server 2007 are the same. In Active Directory, the domain hierarchy is organized using a number of constructs to make administration simpler and more logical. These logical constructs, which are described in the following sections, allow you to define and group resources so that they can be located and administered by name rather than by physical location.

Objects

An *object* is the basic unit in Active Directory. It is a distinct named set of attributes that represents something concrete, such as a user, printer, computer, or application. *Attributes* are the characteristics of the object; for example, a computer is an object and its attributes include its name and location. A user is also an object. In Exchange, a user's attributes include the user's first name, last name, and email address. User attributes also include Exchange-related features, such as whether the object can receive email, the formatting of email it receives, and the location where it can receive email.

Organizational Units

An *organizational unit (OU)* is a container in which you can place objects such as user accounts, groups, computers, printers, applications, file shares, and other organizational units. You can use organizational units to hold groups of objects, such as users and printers, and you can assign specific permissions to them. An organizational unit cannot contain objects from other domains, and it is the smallest unit to which you can assign or delegate administrative authority. Organizational units are provided strictly for administrative purposes and convenience. They are transparent to the end user but can be extremely useful to an administrator when segmenting users and computers within an organization.

You can use organizational units to create containers within a domain that represents the hierarchical and logical structures within your organization. This enables you to manage how accounts and resources are configured and used.

You can also use organizational units to create departmental or geographical boundaries. In addition, you can use them to delegate administrative authority over particular tasks to particular users. For instance, you can create an OU for all your printers and then assign full control over the printers to your printer administrator.

Domains

A *domain* is a group of computers and other resources that are part of a network and share a common directory database. A domain is organized in levels and is administered as a unit with common rules and procedures. All objects and organizational units exist within a domain. Also, all domains are part of a forest, which is a collection of domains.

You create a domain by installing the first domain controller inside it. In Windows Server 2008, a domain controller is created first by installing the Active Directory Domain Services role. Once the role has been installed, you can use the Active Directory Domain Services Installation Wizard to install Active Directory. To install Active Directory on the first server on a network, that server must have access to a server running as a *Domain Name System* (*DNS*) server. If it does not, the installation wizard will install and configure the DNS service for you.

A domain can exist in one of five possible domain functional levels, as outlined in the following list:

- *Windows 2000 mixed*: The default domain functional level for all new domains. It allows for Windows NT 4.0 backup domain controllers (BDCs), Windows 2000 Server domain controllers, and Windows Server 2003 domain controllers. Local and global groups are supported, but universal groups are not. Global catalog servers are supported. Exchange Server 2007 cannot be installed into a domain with this functional level.

- *Windows 2000 native*: The minimum domain functional level at which universal groups become available, along with several other Active Directory features. It allows for Windows 2000 Server, Windows Server 2003, and Windows Server 2008 domain controllers. Exchange Server 2007 can be installed in a domain with this functional level, but some advanced cross-forest features are not available.

- *Windows Server 2003 interim*: This supports only Windows NT 4.0 and Windows Server 2003 domain controllers. This mode is only used when you upgrade domain controllers in Windows NT 4.0 domains to Windows Server 2003 domain controllers. The domains in a forest are raised to this functional level; the forest level has been increased to interim.

- *Windows Server 2003*: This level provides all Windows Server 2003 features and functionality such as domain rename. This allows for only Windows Server 2003 and higher domain controllers. All cross-forest Exchange Server 2007 features are supported at this functional level.

- *Windows Server 2008*: This provides the highest level of features, such as fine-grained password policies. This level allows for only Windows Server 2008 and higher domain controllers.
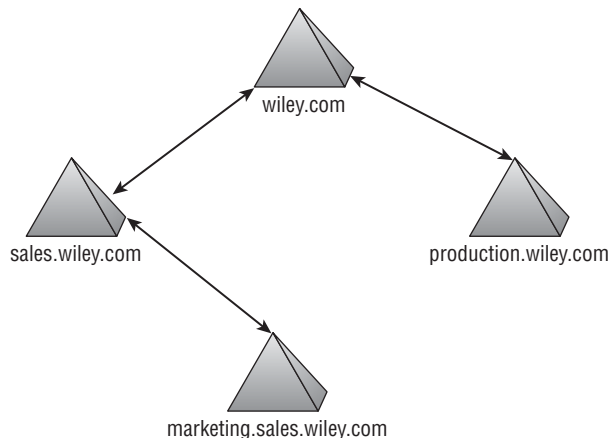
> **WARNING**   The move from a lower functional level to a higher one is irreversible, so be sure that all domain controllers running previous versions of Windows Server have been retired or upgraded before you change the functional level.

Domain Trees

A *domain tree* is a hierarchical arrangement of one or more Windows Active Directory domains that share a common namespace. DNS domain names represent the tree structure. The first domain in a tree is called the *root domain*. For example, a company named Wiley Publishing (that has the Internet domain name wiley.com) might use the root domain wiley.com in its primary domain tree. Additional domains in the tree under the root domain are called *child domains*. For example, the domain sales.wiley.com would be a child domain of the wiley.com domain. Figure 1.1 shows an example of a domain tree.

F I G U R E 1.1   A domain tree is a hierarchical grouping of one or more domains.
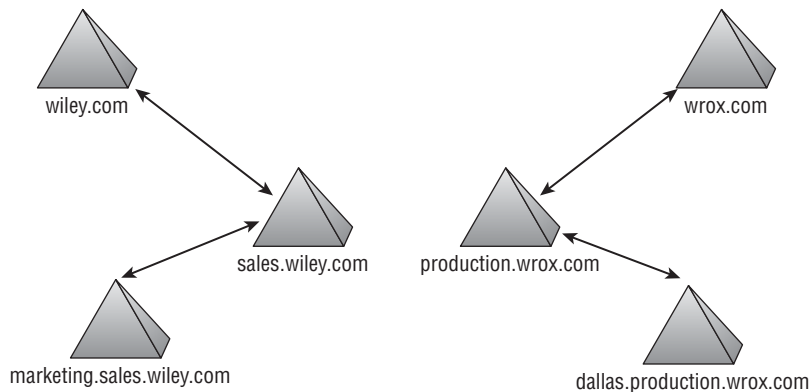
Domains establish trust relationships with one another that allow objects in a trusted domain to access resources in a trusting domain. Since Windows Server 2003, Active Directory has supported transitive, two-way trusts between domains. When a child domain is created, a trust relationship is automatically configured between that child domain and the parent domain. This is a two-way trust, meaning that resource access requests can flow from either domain to the other. The trust is also transitive, meaning that any domains trusted by one domain are automatically trusted by the other domain. For example, in Figure 1.1, consider the three domains named wiley.com, sales.wiley.com, and marketing.sales.wiley.com. When sales.wiley.com was created as a child domain of wiley.com, a two-way trust was formed between the two. When marketing.sales.wiley.com was created as a child of sales.wiley.com, another two-way trust was formed between those two domains. Though no explicit trust relationship was ever defined directly between the marketing.sales.wiley.com and wiley.com domains, the two domains trust each other anyway because of the transitive nature of trust relationships.

Domain Forests

A *domain forest* is a group of one or more domain trees that do not form a contiguous namespace but might share a common schema and global catalog. There is always at least one forest on the network, and it is created when the first Active Directory–enabled computer (domain controller) on a network is installed. This first domain in a forest is called the *forest root domain*, and it is special because it is really the basis for naming the entire forest. It cannot be removed from the forest without removing the entire forest. Finally, no other domain can ever be created above the forest root domain in the forest domain hierarchy. Figure 1.2 shows an example of a domain forest with multiple domain trees.

FIGURE 1.2   A domain forest consists of one or more domain trees.

wiley.com

sales.wiley.com

marketing.sales.wiley.com

wrox.com

production.wrox.com

dallas.production.wrox.com

A forest defines the outermost boundary of Active Directory; the directory cannot be larger than the forest. You can create multiple forests and then create trust relationships between specific domains in those forests; this would let you grant access to resources and accounts that are outside a particular forest. However, an Exchange organization cannot span multiple forests.

Physical Components

The physical side of Active Directory is primarily represented by domain controllers and sites. These enable organizations to optimize replication traffic across their networks and to assist client workstations in finding the closest domain controller to validate logon credentials.

Domain Controllers

Every domain must have at least one *domain controller*, a computer running Windows Server that validates user network access and manages Active Directory. During the process of configuring Active Directory, you have the option of creating a new domain or joining an existing domain. If you create a new domain, you also have the option of creating or joining an existing domain tree or forest. A domain controller stores a complete copy of all Active Directory information for that domain, manages changes to that information, and replicates those changes to other domain controllers in the same domain. Schema and infrastructure configuration information are replicated between all domain controllers in a forest.

> **NOTE** In Windows Server NT 4.0, a distinction was drawn between primary and backup domain controllers. This distinctions meant that the primary domain controller was the only server that changes to objects in the directory. Since Windows 2000 Server was released, all domain controllers are considered peers, and each holds a writable copy of Active Directory. As you talk to people in the IT industry, you may hear them improperly refer to additional Active Directory domain controllers as backup domain controllers (BDCs). Usually when they say this they mean that an additional domain controller is required. Introduced with Windows Server 2008 was a new option similar to BDCs called read-only domain controllers (RODC).

Global Catalog

In a single-domain environment, users can rely on Active Directory to provide all of the necessary information about the resources on the network. In a multiple-domain environment, however, users often need to access resources outside their domain—resources that might be more difficult to find. For this, a *global catalog* holds information about all objects in a forest. The global catalog enables users and applications to find objects in an Active Directory domain tree if the user or application knows one or more attributes of the target object.

Through the replication process, Active Directory automatically generates the contents of the global catalog from the domain controllers in the directory. The global catalog holds a partial replica of Active Directory. Even though every object is listed in the global catalog, only a limited set of attributes for those objects is replicated in it. The attributes listed for each object in the global catalog are defined in the schema. A base set of attributes is replicated to the global catalog, but you can specify additional attributes to meet the needs of your organization.

NOTE — By default, the entire forest has only one global catalog, and that is the first domain controller installed in the first domain of the first tree. All others must be configured manually. It is recommended to always add a second global catalog for backup and load balancing. Furthermore, each domain should have at least one global catalog to provide for more efficient Active Directory searches and network logons.

Active Directory Sites

An Active Directory site is a group of computers that exists on one or more IP subnets. Computers within a site should be connected by a fast, reliable network connection. Using *Active Directory sites* helps maximize network efficiency and provide fault tolerance. DNS also uses Windows sites to help clients find the closest domain controller to validate logon credentials.

NOTE — Exchange Server 2007 makes extensive use of Active Directory information on global catalog servers. For efficient communication, Exchange Server 2007 requires a local global catalog server.

Sites are created and configured using the Active Directory Sites and Services tool. No direct relationship exists between Windows domains and sites, so a single domain can span multiple sites, and a single site can span multiple domains.

One of the biggest mistakes administrators make is not defining all of the IP subnets and assigning them to the appropriate sites. Failing to do this can block the installation of Exchange Server 2007, causing mail-flow issues and client-connectivity problems.

NOTE — Before deploying Exchange Server 2007, review the Active Directory site configuration to ensure all IP subnets are defined and assigned properly.

Schema

A *schema* represents the structure of a database system—the tables and fields in that database and how the tables and fields are related to one another. The Active Directory information is also represented by a schema. All objects that can be stored in Active Directory are defined in the schema.

Installing Active Directory on the first domain controller in a network creates a schema that contains definitions of commonly used objects and attributes. The schema also defines objects and attributes that Active Directory uses internally. When Exchange Server 2007 is installed, Exchange setup extends the schema to support information that Exchange needs. Updates to the schema require replication of the schema across the forest and to all domain controllers in the forest. Additional information about how Exchange updates the schema will be presented in Chapter 2, "Installing Exchange Server 2007."

Active Directory Partitions, Masters, and Replication

The information contained within Active Directory is not all stored in a single location, or partition in this case. Actually, five Active Directory partitions contain different pieces of information about the Active Directory forest and domains. Because each partition type contains different information, the domain controllers to which each partition type is replicated within the forest are also different. We'll briefly cover these directory partitions in the following sections.

Domain Partition

The domain partition contains all the objects that you as an administrator are used to working with on a daily basis. These objects include items such as user accounts, computer accounts, and groups. The contents of the domain partition thus are specific to each individual domain within a forest and therefore are replicated only to the domain controllers in that specific domain.

Configuration Partition

The configuration partition contains all the configuration information about the forest, including information about Active Directory and AD-integrated applications such as Exchange Server. As such, the configuration partition is replicated to every domain controller in the entire forest. Applications benefit from storing their configuration data in the configuration partition because no additional work or configuration is needed to ensure that configuration information is available forest-wide.

Schema Partition

The schema partition, true to its name, is the housing location for the information that defines what objects exist within that Active Directory forest. Each object also has multiple attributes that can be configured, and thus they are also defined in the schema. The schema partition, being so critical to Active Directory, is also replicated to every domain controller in the forest. Unlike other data in Active Directory, only one copy of the schema partition is writable; that is to say that only one domain controller can make changes to the schema. This special domain controller is known as the *schema master.* We'll discuss the "roles" that domain controllers hold, including that of the schema master, later in this chapter.

Application Partition

Application partitions were introduced in Windows Server 2003 and were designed for holding data that is specific to an application. By default, no application partitions are created in a fresh installation of Active Directory; however, some are usually created to store Active Directory–integrated DNS zones. Application partitions are not limited to being replicated to only a single domain or the entire forest—replication can occur with any domain controller in the forest, spanning multiple domains.

Global Catalog Partition

The global catalog partition is a special type of Active Directory partition that is replicated to configured domain controllers across the entire forest. The global catalog is a read-only partial representation of all objects in a forest. It is used to locate information about objects from any domain in the forest without having to know in which domain the object is located. Exchange heavily relies on the global catalog server to resolve recipient email addresses.

Active Directory Masters

Within each Active Directory forest, five unique "roles" exist (or more properly, *operations masters*) that reside on certain domain controllers. Active Directory uses a multimaster replication system, which means that all domain controllers are equal. Well, mostly equal. Certain tasks do not lend themselves well to having multiple domain controllers performing them (especially at the same time), so the operations master roles exist. Active Directory has five operations master roles, and initially all five exist on the first domain controller installed in a new forest. You can, and should, move roles around as additional domain controllers are joined to the forest and as subsequent domains are created within the forest.

The operations master roles in Active Directory are as follows:

- *Domain naming master*: Only one domain naming master role exists within an entire forest. The domain controller that holds this role is responsible for creating new domains in the forest and also for removing domains from the forest. These tasks cannot normally be performed if the domain controller holding this role is unavailable.

- *Schema master*: There is only one schema master role within a forest. As we discussed briefly, any changes that need to be made to the schema of the forest must be made on the schema master. Once the changes are made on the domain controller holding this role, they are replicated to the rest of the domain controllers in the forest. A failure of the schema master will prevent any schema modifications from being made in that forest. Exchange Server 2007, specifically, requires schema modifications and thus will fail to install if the schema master cannot be contacted.

- *Infrastructure master*: An infrastructure master role exists in each domain in the forest and is responsible for updating changes made to user account names and group memberships. The domain controller holding this role in the domain maintains the up-to-date copy of this information as it is changed and then replicates it to the other domain controllers in the domain.

- *PDC emulator master*: One primary domain controller (PDC) emulator master role exists in each domain in the forest. The PDC emulator master is required to provide backward interoperability with Windows NT 4.0 backup domain controllers (BDCs). In a mixed environment, the PDC emulator master processes all password changes in the domain. Additionally, failed authentication attempts are forwarded to the PDC emulator to be retried, accounting for changes that might have occurred to the password. The PDC emulator master also typically functions as the Network Time Protocol (NTP) source for the domain and is usually configured to take time input from a trusted internal (that is, atomic or satellite clock) or external NTP time source.

- *RID master*: One RID master role exists in each domain in the forest and is responsible for issuing blocks of relative identifiers (RIDs) to other domain controllers in the domain. This block of RIDs is known as the *RID pool*. When a domain controller runs low on RIDs in its RID pool, it makes a request to the RID master for another block of RIDs for its usage. Each object that exists within a domain has a unique security identifier (SID). This SID is composed of two parts: a domain RID (common throughout the domain) and a unique RID from the RID pool. These are combined to create a globally unique (within the forest) SID for that object. When the pool of RIDs has been exhausted on a domain controller, it will be unable to create new objects in the domain. Exchange Server 2007 creates several security principals during its installation and thus requires the usage of some RIDs from the RID pool of a domain controller.

Replication

Although I've mentioned replication in Active Directory several times, I have not yet properly discussed it. I'll remedy that situation now before moving into the next section of this chapter.

Replication is the process by which all domain controllers in a domain or forest pass changes to other domain controllers and thus update their copies of the specific Active Directory partitions they hold as they themselves receive replication updates from other domain controllers. Because changes occur almost constantly across multiple domain controllers within a forest, the replication used for Active Directory is referred to as *loosely consistent*, meaning that not every domain controller in the forest with a certain partition will have the same information at any time. However, over time, *convergence* occurs as all domain controllers receive and pass replication updates and the partitions that they hold become closer to matching exactly. In a production environment with multiple domain controllers, complete convergence is almost impossible to achieve, but that rarely poses a problem. Administrators with the appropriate permissions can always manually trigger replication to be performed between domain controllers, so important changes can be forced to replicate if normal replication schedules are not appropriate at the time, which is typically a problem only when dealing with *intersite replication*.

Given that Active Directory uses sites to map the Active Directory network to that of the physical network, replication thus occurs differently between sites (intersite replication) than it does between domain controllers in the same site (intrasite replication). Intersite replication is designed to have the minimum possible impact on the typically slower wide area network (WAN) links that commonly separate the physical locations that Active Directory sites represent. As such, the replication traffic is highly compressed and also occurs on a schedule that is configured on the *site link* object that is created to logically connect two Active Directory sites. Thus, changes made on a domain controller in Site A will not be sent to a domain controller in Site B until the next scheduled replication time based on the replication interval and allowable replication times that were configured. Conversely, intrasite replication occurs almost immediately after a change has been made to some bit of Active Directory information. The domain controller that the change is made on will wait 15 seconds (to account for any additional changes) and then will begin replicating its changes to

the other domain controllers within that site. After replication has occurred with the first replication partner that domain controller has, it will wait three seconds and then commence replication with its next replication partner, and so forth, until the original domain controller has replicated with all replication partners within that site.

Replication latency occurs when a change made on one domain controller has not been replicated to another domain controller, either in the same site or in a different site. Obviously, the replication latency within a site should always be much lower than that between sites, but should replication problems arise between domain controllers, latency can even exist within a site. On the surface, replication latency is not completely desirable, and it must be dealt with accordingly when using a distributed multimaster replication environment like Active Directory uses. Faster, higher-quality (or cheaper) WAN links will lend themselves to configuration replication occurring more frequently than slower, less reliable (or more expensive) WAN links. The price to be paid for lower replication latency in an Active Directory environment that is otherwise healthy and functioning properly is the cost of pushing more data over these typically congested and high-cost WAN links. The administrator who is ultimately responsible for managing Active Directory across your organization will have to determine what is best to meet their specific needs.

> **NOTE** To learn more about Active Directory, start by checking out the Windows Server 2008 product documentation. It provides an overview of the technology and illustrates many of the benefits of using Active Directory.

Active Directory and Exchange Server 2007

In versions prior to Exchange 2000 Server, Exchange maintained a directory of its own through a service known as the Directory Service. The Directory Service maintained a copy of the directory in a database file on each Exchange server and took care of replicating changes in the directory to other Exchange servers.

Since Exchange 2000 Server, Exchange has been totally reliant on Active Directory to provide its directory services. This reliance causes a shift in the way that the Exchange directory is maintained. The section "Forests" examines the effects that boundaries of a forest place on Exchange. Then the section "Domain Name Service (DNS)" looks at the interaction of DNS in an Exchange organization.

Forests

By default, the global catalog shows objects only within a single Active Directory forest, so an Exchange organization must be within the boundaries of a forest. This is different from earlier versions of Windows NT and Exchange 5.5. In previous versions, an Exchange organization could span domains that did not trust one another because Exchange 5.5 did not rely so much on the underlying security structure of Windows NT. With Active Directory and Exchange Server 2007, the security structure is integrated, which means a single

Exchange organization cannot span multiple forests but can span multiple domains within a single forest.

Domain Name Service (DNS)

For Active Directory and Exchange Server 2007 to function, DNS must be properly functioning in your organization. Outlook Web Access, SMTP connectivity, and Internet connectivity all rely on DNS.

Active Directory is often called a *namespace*, which is similar to the directory service in earlier versions of Exchange and means any bounded area in which a given name can be resolved. The DNS name creates a namespace for a tree or forest, such as wiley.com. All child domains of wileypublishing.com, such as sales.wiley.com, share the root namespace. In Exchange Server 2007, Active Directory forms a namespace in which the name of an object in the directory can be resolved to the object. All domains that have a common root domain form a *contiguous namespace*. This means the domain name of a child domain is the child domain name appended to the name of the parent domain.

In Active Directory, a DNS domain name such as exchangeexchange.local does not directly affect the email addresses for Exchange users created in that domain. Although a user's logon name might be user@exchangeexchange.local, you can control how email addresses are generated using email-address-generation policies in the Exchange Management Console.

Active Directory Partitions

Although you've examined briefly already how Exchange Server 2007 uses the different Active Directory partitions, you'll dig a bit deeper in this section. Recall that there can be only one Exchange Server organization within an entire forest. Specific examples of how Exchange Server 2007 uses these Active Directory partitions include the following:

- The configuration partition stores all configuration information about the Exchange organization. This information includes items such as recipient policies, address lists, and Exchange settings. The configuration partition is replicated to every domain controller in the forest; therefore, this critical Exchange configuration information is available to every domain user irrespective of which domain holds their user account.

- The domain partition stores information about the basic blocks of Exchange Server: its recipient objects. Recipient objects include the users, contacts, and groups that have configured email addresses. We'll go into great depth about configuring and managing recipients in Chapter 5, "Working with Recipients, Groups, and Mailboxes."

- The schema partition is modified by the Exchange Server 2007 setup routine to add attributes to existing objects, such as users and groups. Additionally, the schema is extended to include Exchange Server–specific objects that are required for Exchange Server to function properly. We'll cover modifying the schema to support the installation of Exchange Server 2007 in Chapter 2, "Installing Exchange Server 2007."

- The global catalog partition received many new items of information as a result of the installation of Exchange Server 2007 in a forest. Exchange uses the global catalog to generate address lists for usage by Exchange recipients, and Exchange Server also uses it to locate a recipient to aid in the delivery of mail items to that recipient. Exchange Server automatically generates the global address list (GAL) from all recipients listed in the global catalog.

Message Flow

In previous versions of Exchange Server, a complex link-state routing algorithm was used to route messages between geographically separated Exchange servers. Exchange used routing groups that were connected with routing-group connectors to perform this routing. With the elimination of routing groups and link-state routing in Exchange Server 2007, all Exchange message routing is performed by Hub Transport servers using the Active Directory sites and site links that service Active Directory itself. As such, message routing (both within the same site and across site links) is significantly less complex in Exchange Server 2007.

> We will cover Exchange Server 2007 routing in Chapter 3, "Configuring the Mailbox and Hub Transport Roles."

Within each Active Directory site that contains a Mailbox server (or Unified Messaging server), you must have at least one Hub Transport server. The Hub Transport server is responsible for routing all messages within a site and between connected sites. Even a message that is sent from a recipient on Server A to another recipient on Server A must first cross through a Hub Transport server for delivery—a big change in message routing from Exchange Server 2003. When messages must be routed between sites, the Hub Transport server in the originating site determines the best route available at that time to the destination server and routes the message accordingly.

Message routing between sites occurs as detailed here:

1. The sending user submits the message to their mailbox on the Mailbox server.

2. The Mailbox server notifies a Hub Transport server in its Active Directory site that it has a message awaiting pickup.

3. A Hub Transport server in the same Active Directory site as the originating Mailbox server picks up (retrieves) the message from the Mailbox server.

4. The Hub Transport server performs a query against Active Directory to determine what Mailbox server the recipient of the message is on.

5. The Hub Transport server then computes the lowest-cost route to the site containing the destination Mailbox server based on the site link costs configured on site links between the sites.

6. The Hub Transport server in the originating Active Directory site then sends the message along the lower-cost route it has computed.

7. If multiple Active Directory sites must be crossed, the message is delivered to a Hub Transport server along the path and then passed along to a Hub Transport server in the destination site.

8. If there are no operating Hub Transport servers in the destination site, the message will be queued on a Hub Transport server in the site closest to the one where the destination Mailbox server resides. The message will not be delivered until a Hub Transport server in the destination site is available to deliver it.

9. When the message reaches the Hub Transport in the destination site, that Hub Transport server assumes responsibility to deliver the message, and the message is sent to the appropriate destination Mailbox server.

What's New in Exchange Server 2007?

With any new release of an established product like Exchange Server, Microsoft includes new (and improved) features that benefit both the administrative side of the product and the end-user experience. I'll briefly highlight some of the key features that are new or improved in Exchange Server 2007 (although this list is certainly not all-inclusive):

- *Exchange Management Console*: The first, and most striking, change that many administrators with Exchange experience will notice is that the familiar Exchange System Manager is gone and has been replaced by the completely redesigned Exchange Management Console (EMC). By examining the ways administrators worked and the tasks they needed to perform, Microsoft designed the EMC to be as intuitive and workflow-oriented as possible. The EMC also takes advantage of the improvements in Microsoft Management Console 3.0. We will spend a good deal of our time together in this book working with the EMC.

- *Exchange Management Shell*: Another dramatic change from an administrative standpoint is the Exchange Management Shell, which is a new command-line shell and scripting environment for Exchange administrators. Any action that can be carried out in the EMC can be performed just as easily in the Exchange Management Shell, and many actions that an Exchange administrator will perform can be performed only from within the Exchange Management Shell. You'll see as you work with Exchange Server 2007 that almost every configuration action you perform in the EMC will present you with the corresponding Exchange Management Shell code that is actually being used to carry out those changes.

- *64-bit*: Exchange Server 2007 is the first messaging platform to utilize the benefits of 64-bit hardware and operating systems fully. In fact, Exchange Server 2007 is available for production use only in 64-bit versions. The amount of RAM available to be used efficiently in 64-bit environments is significantly higher than in 32-bit environments, thus allowing for more mailboxes and storage groups on a single Exchange server.

- *Active Directory (AD) site–based routing*: No longer do you need to plan, implement, and manage an Exchange-specific routing environment with routing groups. Exchange Server 2007 is AD site–aware and will use the existing Active Directory sites configuration to perform routing and to select which Exchange servers it should directly communicate with. This change will allow a closer alignment of the physical network topology with the Exchange routing topology.

- *Server roles*: Gone are the days of every Exchange installation being the same as every other installation. Also gone are the days of a single check box being the determining factor in the role played by an Exchange server. Exchange Server 2007 now allows—in fact, it demands—you to deploy it in one or more of several available roles. The familiar back-end server of old is now referred to as a Mailbox server, although it can certainly still host public folders. The closest role to that of the old front-end server would be that of the Client Access server. You'll examine all the roles, uses, benefits, and limitations of Exchange Server in detail in Chapters 2 and 3.

- *Unified messaging*: Once a popular, complex, and costly third-party add-on for Exchange, unified messaging is now available within Exchange Server 2007 by deploying the Unified Messaging role and using Exchange Server 2007 Enterprise client access licenses (CALs). Unified messaging is outside the scope of the 70-236 exam, so I will not be discussing it in this book.

- *Highly available*: In the past, if you wanted highly available Exchange servers, you had two choices from Microsoft: active/passive clusters or active/active clusters. Both were certainly suitable but complex and costly—a reality that prevented many smaller organizations from providing a highly available Exchange solution. Additionally, there were third-party applications that promised various high-availability solutions for Exchange Server, and many of them were very good products. Seeing the need to revamp the high-availability solutions offered in Exchange and wanting to take advantage of new clustering capabilities, Microsoft introduced three new high-availability solutions in Exchange Server 2007: local continuous replication (LCR), cluster continuous replication (CCR), and standby continuous replication (SCR). Clustering using active/passive nodes has been improved and renamed to single copy clustering (SCC), while support for active/active clustering has been eliminated entirely. You will examine high availability for Exchange Server 2007 in Chapter 8, "Configuring Highly Available Exchange Server Solutions."

- *Compliance and message management*: As email continues to grow and evolve as the number-one means of business-critical communication, the need to manage and enforce certain policies on email content and usage also grows. Exchange Server 2007 presents several novel, and quite useful, methods that allow organizations to control the growth of the messaging stores and also to monitor and control the usage of email, thus protecting the organization from legal or other troubles. You'll examine compliance and message management in Exchange Server 2007 in Chapter 7, "Configuring Security, Compliance, and Policies."

▪ *Antivirus and antispam controls*: The Edge Transport role, one of the new Exchange Server 2007 server roles, is responsible for preventing spam messages from entering your Exchange organization. The intelligent message filter (IMF) has been removed from the Exchange servers that host mailboxes and public folders or that handle client access requests and moved into the Edge Transport role, which is designed to operate in a demilitarized zone (DMZ) network if desired. Additionally, Sybari's Antigen antivirus product is now a Microsoft product known as Forefront Security for Exchange Server. Forefront is a complete Exchange-aware antivirus application that can be used on the Edge Transport server as a network edge scanner and also on the Hub Transport server to scan messages traversing your internal network. You'll examine antivirus and antispam issues in more detail in Chapter 7.

What's No Longer Supported in Exchange Server 2007?

In any new release of a software product, discontinued or de-emphasized features are inevitable. Such is the case with Exchange Server 2007, although some of these items might surprise experienced Exchange administrators. The items that follow in no way represent every change that has occurred in Exchange Server 2007, but they do represent some of the most interesting ones.

Features That Have Been Removed or Replaced

The following key features and functionality have been removed from Exchange Server 2007:

▪ *Routing groups*: Link-state routing is no longer used in Exchange Server 2007 and has been replaced by Active Directory site–based routing. This places further importance on the proper planning and design of the Active Directory forest into which Exchange Server 2007 will be installed, but it reduces the overall amount of planning and administration required to maintain an Exchange organization. Now all routing (both AD and Exchange) is controlled and configured from a single location—the Active Directory Sites and Services console—thus providing consistent, predictable results that can be controlled as your physical network dictates. You'll examine Active Directory more as it relates to the installation of Exchange Server 2007 in Chapter 2.

▪ *Administrative groups*: Administrative groups, which were previously used in Exchange Server to control administrative access to groups of servers, have been replaced by the Exchange Server 2007 split permissions model that emphasizes using universal security groups. We'll cover administrative roles more in Chapter 3.

- *Exchange management via Active Directory Users and Computers*: Management of all recipient objects (discussed more in Chapter 5) is now performed via the Exchange Management Console and Exchange Management Shell. Management of Exchange recipients has been integrated in the Active Directory Users and Computers (ADUC) console in the previous two versions of Exchange Server, but Exchange administrators who've worked with Exchange Server 5.5 will recall this method of management very well.

- *Streaming database*: The streaming database (`*.stm`), first introduced in Exchange 2000 Server, has been removed in Exchange Server 2007.

- *Recipient Update Service*: The Recipient Update Service (RUS) has been removed from Exchange Server 2007. To find out more about how to work with recipients, you will examine email address generation more in Chapter 5.

- *Exchange 5.5 interaction*: Exchange Server 2007 does not interoperate with the Active Directory Connector (ADC) or Site Replication Service (SRS) as in the previous two versions of Exchange. As a result, you cannot directly migrate from Exchange Server 5.5 to Exchange Server 2007. We'll discuss migration briefly in Chapter 2.

- *Network News Transfer Protocol (NNTP)*: This has been removed completely. You'll need to use Exchange Server 2003 or Exchange 2000 Server to provide this protocol to clients.

- *X.400 message transfer agent*: This has been removed completely. You'll need to use Exchange Server 2003 or Exchange 2000 Server if your organization needs this message transfer agent protocol.

- *Novell GroupWise connector*: This has been removed completely. You'll need to use Exchange Server 2003 or Exchange 2000 Server to provide this connector.

- *Lotus Notes connector*: This is no longer available, but Microsoft has provided migration and coexistence tools for Exchange Server 2007 by using the Microsoft Transporter Suite tools.

- *Active/active clustering*: This is no longer supported. You'll need to implement either an active/passive SCC model or consider using the new high-availability features provided by CCR. You'll spend all of Chapter 10, "Disaster Recovery Operations for Exchange Server," looking at highly available Exchange Server 2007 implementations.

- *IMAP4 access to public folders*: You'll need to retain Exchange Server 2003 or Exchange 2000 Server to provide IMAP4 access to public folders to clients.

- *Exchange WebDAV extensions*: Exchange WebDAV has been replaced by the Exchange Web Services.

Features That Have Been De-emphasized

The following key features and functionality have been de-emphasized in Exchange Server 2007:

- *Public folders*: Public folders are no longer required in a clean installation of Exchange Server 2007. In previous versions of Exchange Server, public folders contained critical system data such as the Offline Address Book (OAB) and free/busy calendaring data. This is no longer the case, because no system data is stored in public folders in a pure Exchange 2007 environment. Public folders, however, are still supported in Exchange Server 2007, although Microsoft recommends moving to SharePoint Portal Server or another product for those items that previously were stored in public folders. It's expected that public folders (which were initially advertised as not being supported in Exchange Server 2007) will not be supported in a future release of Exchange Server. Microsoft's official support policy says that they will support Public Folders in Exchange 2007 until 2016. I'll cover public folders in Exchange Server 2007 in Chapter 6, "Configuring and Managing Public Folders."

- *Exchange Server 2003 virus scanning API (VSAPI)*: Although Exchange Server 2007 still supports the VSAPI, its role is being de-emphasized because Microsoft has started to integrate antivirus controls at the transport layer. We'll cover antivirus controls in Exchange Server 2007 in Chapter 7.

> **NOTE** You can look at the entire list of new and removed features in Exchange Server 2007 by visiting the TechNet website at `http://technet.microsoft.com/en-us/library/aa996018.aspx`.

Preparing the Infrastructure for Exchange Server 2007 Deployment

If you are planning to deploy Exchange Server 2007, you need to make sure that all of the requirements are met. We have covered many of these requirements over the last few sections in this chapter. These requirements and limitations encompass much of what you need to know to be successful in preparing to install Exchange Server 2007, and to succeed on the exam.

- The schema master must be running Windows Server 2003 Service Pack 1 or later.

- A global catalog running Windows Server 2003 Service Pack 1 or later must be available in each Active Directory site into which Exchange will be deployed. This cannot be a Windows Server 2008 read-only domain controller.

- The Active Directory forest must be set to at least the Windows Server 2000 functional level; however, if cross-forest free/busy information will be shared or forest-to-forest delegation is used, the forest must be set to at least the Windows Server 2003 functional level.

- The Active Directory domain that Exchange will be installed into or that contains Exchange recipients should use at least the Windows 2000 Server native domain functional level.

- All Exchange Server 5.5 computers must be removed from the domain and the Exchange organization must be set to Exchange 2000 native mode.

- Exchange Server 2007 and Exchange Server 2007 Service Pack 1 do not support renaming of the Active Directory domain. If you are considering performing a domain rename, you must do so before deploying Exchange Server 2007.

- Any Exchange 2000 servers and machines with the Exchange 2000 management tools installed in the environment must have the latest post–Service Pack 3 rollup applied.

- Any Exchange Server 2003 computers and machines with the Exchange 2003 management tools installed must have Service Pack 2 applied.

- Be sure that you have verified that your backup, file-level antivirus, third-party mobile messaging, and alert monitoring systems will support Exchange Server 2007.

⊕ Real World Scenario

Planning an Active Directory Deployment

If you are planning a completely new Active Directory deployment for your organization, then you should be certain to place domain controllers and global catalog servers in locations that make sense for how your company is organized and how it operates. When planning how and where to locate these key servers in your Active Directory environment, there is no absolute answer that works for all scenarios. The saying "the more, the better" is not necessarily true, especially if replication over slow WAN links becomes too much for those links to handle. Conversely, saying "less is more" is almost always untrue when it comes to implementing a solid Active Directory infrastructure. Remember, this will be the foundation of your entire network, so you should take however long you need to get it right the first time.

These are a few general guidelines you should keep in mind as you're working in different scenarios:

- Every domain in the Active Directory forest should have at least two domain controllers. This is for both client load balancing and disaster recovery in case one domain controller should fail.

- You should place additional domain controllers in domains as organizational structures (such as physical location or client groupings) dictate.

- You should be aware that additional domain controllers will cause additional replication traffic, which can be problematic for intersite replication across slow WAN links.

- Every Active Directory site must have at least one domain controller and that one domain controller must be configured as a global catalog if Exchange servers or users are in that site.

- If a site has multiple domain controllers, consider using a Bridgehead server for Active Directory replication to other sites.

- Install the right number of Global Catalog servers in each site to support the applications. When using 32-bit domain controllers, the ratio of Global Catalog processor cores to Exchange Mailbox server processor cores is 1 to 4. When using 64-bit Global Catalogs with enough memory to cache the entire Active Directory, the ratio is 1 to 8.

Summary

The better you understand how the Exchange system works, the better you'll be able to plan a viable network and troubleshoot that network when problems occur. This chapter examined three basic aspects of Exchange Server architecture: how Exchange is integrated with Active Directory, how information is stored and organized in Active Directory, and how messages flow within an Exchange organization.

At the top of the Active Directory hierarchy is the domain forest, which represents the outside boundary that any Exchange organization can reach. A domain tree is a hierarchical arrangement of domains that share a common namespace. The first domain in a tree is the root domain. Domains added under this are child domains. Within the domain tree, domains establish trust relationships with one another that allow objects in a trusted domain to access resources in a trusting domain. A domain is a group of computers and other resources that are part of a network and share a common directory database. Each domain contains at least one domain controller. Multiple domain controllers per domain can be used for load balancing and fault tolerance.

When Exchange is installed, many objects, such as users, are enhanced with Exchange-related features. A global catalog holds information about all the objects in a forest. Objects can be grouped into organizational units that allow administrators to manage large groups of similar objects effectively at the same time.

Within Active Directory, five partitions store certain pieces of the total information that makes up Active Directory. These partitions are the domain partition, configuration partition, schema partition, global catalog partition, and application partition(s). There can be multiple application partitions within the forest and domains.

Although Active Directory uses multimaster replication, there are five specific roles that only one domain controller in a forest or domain can hold at any one time. The five roles are the domain naming master (one per forest), schema master (one per forest), infrastructure master (one per domain), PDC emulator master (one per domain) and RID master (one per domain). The failure of a domain controller holding each role will have different effects on the forest and domain. Exchange Server 2007 must contact the domain controller holding the schema master role during setup to modify and extend the schema.

Active Directory is loosely consistent, meaning that not every domain controller in the forest with a certain partition will have the same information at any time. However, over

time, convergence occurs as all domain controllers receive and pass replication updates and the partitions that they hold become closer to matching exactly. In a production environment, complete convergence is almost impossible to achieve, but that rarely poses a problem.

Intersite replication is designed to have the minimum possible impact on the typically slower WAN links that commonly separate the physical locations that Active Directory sites represent. As such, the replication traffic is highly compressed and also occurs on a schedule that is configured on the site link object that is created to connect two Active Directory sites logically. Conversely, intrasite replication occurs almost immediately after a change to some bit of Active Directory information has taken place. The domain controller that the change is made on will wait 15 seconds (to account for any additional changes), and then it will begin replicating its changes to the other domain controllers within that site.

As you prepare to deploy Exchange Server 2007, you must ensure that the schema master for the forest has at least Windows Server 2003 Service Pack 2 applied and that the domain and forest functional level is set to Windows 2000 native at a minimum. Next you will want to check to make sure that at least one global catalog at each site that Exchange Server 2007 will be installed in has Windows Server 2003 Service Pack 1, or later, applied.

Exam Essentials

Understand Active Directory.   Although this book is not trying to prepare you for an exam related to Active Directory design, support, or administration, it is absolutely imperative that you understand how Active Directory is designed and how it functions. With Exchange being completely Active Directory–integrated and –aware, all administrative functions related to users and mailboxes are tied into Active Directory. To that end, ensure that you have a good understanding of both the logical and physical structure of Active Directory. In addition, you should understand the various domain functional levels that are available in Windows Server 2003 and Windows Server 2008 and how they will impact your overall network.

Understand basic message routing.   It is helpful, both in preparing for this exam and in the day-to-day administration of Exchange Server 2007, to understand how messages are routed within the same site and between different sites. All messages are routed through the Hub Transport server, even if the originating and destination recipients reside on the same Mailbox server.

Understand the infrastructure requirements.   Exchange is so tightly integrated with Active Directory that specific requirements need to be met on the versions and placement of domain controllers. Understand the domain and forest functional levels that need to be set to install Exchange Server 2007.

Review Questions

1. You are currently running in the Windows 2000 mixed domain functional level and are considering making the switch to the Windows 2000 native domain functional level. Which of the following would be valid concerns to take into account before making the switch? (Choose all that apply.)

    A. The switch is irreversible.

    B. If you later decide to switch to the Windows 2000 mixed domain functional level, all object configuration will be lost.

    C. Exchange Server 5.5 cannot be run in a Windows 2000 native domain functional level environment.

    D. You must upgrade or retire all Windows NT 4.0 domain controllers.

2. Which of the following statements is true of domains in a single-domain tree?

    A. Domains are not configured with trust relationships by default.

    B. Domains are automatically configured with one-way trust relationships flowing from parent domains to child domains.

    C. Domains are automatically configured with two-way nontransitive trusts.

    D. Domains are automatically configured with two-way transitive trusts.

3. By default, how long will a domain controller wait to initiate replication to its replication partners in the same Active Directory site after a change is made on it?

    A. 3 seconds

    B. 3 minutes

    C. 15 seconds

    D. 15 minutes

4. A hierarchical arrangement of one or more Active Directory domains that share a common namespace is referred to as a _____ .

    A. Windows Server 2003 site

    B. domain site

    C. domain tree

    D. domain forest

5. You have just installed the first Windows Server 2008 server on your network and want to make it a domain controller. How would you do this?

    A. The first Windows Server 2003 server is automatically made a domain controller.

    B. Install Active Directory Domain Services role on the computer.

    C. Install DNS on the computer.

    D. Install the schema on the computer.

6. Which of the following statements about an organizational unit is true?

    A. An organizational unit cannot contain objects from other domains.

    B. An organizational unit can contain objects only from other trusted domains.

    C. An organizational unit can contain objects only from other domains in the same domain tree.

    D. An organizational unit can contain objects only from other domains in the same domain forest.

7. What service is the primary provider of name resolution on a Windows Server 2008 network?

    A. X.400

    B. DNS

    C. WINS

    D. SMTP

8. Messages in Exchange Server 2007 are routed by which server?

    A. The global catalog server

    B. The infrastructure master server

    C. The Hub Transport server

    D. The Mailbox server

9. If Exchange Server 2007 fails to contact a certain operations master role holder during installation, the installation process will fail. Which operations master role is this?

    A. Infrastructure master

    B. Schema master

    C. RID master

    D. Domain naming master

10. Message routing between Exchange Server 2007 Mailbox servers uses what method to determine the best route?

    A. Link-state algorithms

    B. Site link costs

    C. Packet latency

    D. Open-shortest-path-first routing

11. Of the following features available in Exchange Server 2003, which are no longer supported in Exchange Server 2007? (Choose two answers.)

    A. Public folders

    B. The streaming database

    C. Command-line management

    D. Integration with Exchange Server 5.5

12. User account objects are found in which Active Directory partition?

    A. Configuration

    B. Global catalog

    C. Schema

    D. Domain

13. What impact does the failure of the domain controller holding the schema master role have on the normal operations of Active Directory?

    A. Active Directory will cease to function properly until the schema master role has been brought back online.

    B. Active Directory will continue to function normally except that schema modifications cannot be processed until the schema master role has been brought back online.

    C. Active Directory will continue to function normally except that intrasite replication will fail until the schema master role has been brought back online.

    D. Active Directory will continue to function normally except that down-level Windows NT 4.0 BDCs will not be able to interact with the domain of which they are a part.

14. To use universal groups in your Active Directory domain, at what minimum domain functional level must you be running?

    A. Windows Server 2003

    B. Windows 2008

    C. Windows Server 2003 interim

    D. Windows 2000 native

15. Which of the following is the smallest object that other Active Directory objects can be placed within and have authority delegated over them?

    A. Organizational unit

    B. Forest

    C. Domain

    D. Site

16. Which domain controllers in an Active Directory environment maintain a copy of the configuration partition?

    A. Certain domain controllers in all domains

    B. All domain controllers in a single domain

    C. All domain controllers in the forest

    D. Certain domain controllers in the forest

17. *Intersite replication* refers to which of the following?

    A. Replication between domain controllers in the same Active Directory site

    B. Replication between domain controllers in different domains

    C. Replication between domain controllers in different forests

    D. Replication between domain controllers in different Active Directory sites

18. Which Active Directory partition is used to create the Exchange address lists?

    A. Configuration

    B. Global catalog

    C. Schema

    D. Domain

19. If the Hub Transport server in the destination site is unavailable, where will a message in routing be queued up temporarily?

    A. On the Hub Transport server in the source site

    B. On the Mailbox server in the destination site

    C. On the Hub Transport server in the destination site

    D. On the Hub Transport server in the site nearest to the destination site

20. What administrative console is used to configure the link costs that Exchange Server 2007 uses when routing messages?

    A. Exchange System Manager

    B. Active Directory Users and Computers

    C. Active Directory Sites and Services

    D. Active Directory Domains and Trusts

Answers to Review Questions

1. **A, D.** The switch to the Windows 2000 native domain functional level is a one-time, one-way switch and is irreversible. Once you have switched to the Windows 2000 native domain functional level, you will no longer be able to have Windows NT 4.0 domain controllers within the organization.

2. **D.** Windows Server 2003 (along with Windows 2000 Server) and Active Directory support two-way transitive trusts between domains. When a child domain is created, a trust relationship is automatically configured between that child domain and the parent domain. This trust is two-way, meaning that resource access requests can flow from either domain to the other.

3. **C.** The domain controller on which the change is made will wait 15 seconds (to account for any additional changes) and then will begin replicating its changes to the other domain controllers within that site. After replication has occurred with the first replication partner of the domain controller, it will wait 3 seconds and then commence replication with its next replication partner, and so forth, until the original domain controller has replicated with all replication partners within that site.

4. **C.** A domain tree is a hierarchical arrangement of one or more Windows Active Directory domains that share a common namespace. Domain Name Service (DNS) domain names represent the tree structure. The first domain in a tree is called the *root domain*.

5. **B.** To create a Windows Server 2008 domain controller, you first install the Active Directory Domain Services role on it and then run DCPromo. During this process, you have the option of creating a new domain or joining an existing domain. If you create a new domain, you also have the option of creating or joining an existing domain tree or forest.

6. **A.** An organizational unit is a container in which you can place objects such as user accounts, groups, computers, printers, applications, file shares, and other organizational units. An organizational unit cannot contain objects from other domains and is the smallest unit to which you can assign or delegate administrative authority. Organizational units are provided strictly for administrative purposes and convenience.

7. **B.** DNS is the primary provider of name resolution for Windows Server 2003–based networks. In fact, the Windows Server 2003 domain structure is based on DNS structure, and Active Directory requires that DNS be used.

8. **C.** All messages in Exchange Server 2007 are routed to their destination mailbox by the Hub Transport server, even if the message is sent between recipients on the same Exchange Mailbox server.

9. **B.** Any changes that need to be made to the schema of the forest must be made on the schema master. Exchange Server 2007 requires schema modifications and thus will fail to install if the schema master cannot be contacted.

10. **B.** The Hub Transport server, which is responsible for message routing in Exchange Server 2007, computes the lowest-cost route to the site containing the destination Mailbox server based on the site link costs configured on Active Directory site links between the sites.

11. B, D. The streaming database (`*.stm`), first introduced in Exchange 2000 Server, has been removed in Exchange Server 2007. Several other enhancements have been made to storage in Exchange Server 2007. Exchange Server 2007 does not interoperate with the Active Directory Connector (ADC) or Site Replication Service (SRS) as in the previous two versions of Exchange. As a result, you can no longer directly migrate from Exchange Server 5.5 to Exchange Server 2007.

12. D. The domain partition contains all of the objects that you as an administrator are used to working with on a daily basis. These objects include user accounts, computer accounts, and groups. The contents of the domain partition thus are specific to each individual domain within a forest and therefore are replicated to the domain controllers in that specific domain only.

13. B. A failure of the schema master will prevent only schema modifications from being made in that forest.

14. D. The Windows 2000 native domain functional level is the minimum domain functional level at which universal groups become available, along with several other Active Directory features; it allows for Windows 2000 Server and Windows Server 2003 domain controllers only.

15. A. The organizational unit (OU) is a container in which you can place objects such as user accounts, groups, computers, printers, applications, file shares, and other organizational units. An organizational unit cannot contain objects from other domains and is the smallest unit to which you can assign or delegate administrative authority.

16. C. The configuration partition contains all the configuration information about the forest, including information about Active Directory and AD-integrated applications such as Exchange Server. As such, the configuration partition is replicated to every domain controller in the entire forest.

17. D. Intersite replication occurs between domain controllers in different Active Directory sites. Intrasite replication occurs between domain controllers in the same Active Directory site. Sites can span domains, and domains can span sites; thus, no direct relationship must exist between the two. Forests do not replicate.

18. B. Exchange uses a global catalog to generate address lists for usage by Exchange recipients and also uses it to locate a recipient to aid in delivering mail items to that recipient. The global address list (GAL) is automatically generated by Exchange Server from all recipients listed in the global catalog.

19. D. If there are no operating Hub Transport servers in the destination site, the message will be queued on a Hub Transport server in the site closest to the one where the destination Mailbox server resides. The message will not be delivered until a Hub Transport server in the destination site is available to deliver it.

20. C. Sites (and site link costs) are created and configured using the Active Directory Sites and Services tool.